Methodologies for Bandwidth Allocation, Transmission Scheduling, and Congestion Avoidance in Broadband ATM Networks

K. Sriram
Room 3H-607, AT&T Bell Laboratories, Holmdel, N.J. 07733

ABSTRACT: The call types supported in high-speed packet networks vary widely in their bandwidth requirements and tolerance to message delay and loss. In this paper, we classify various traffic sources which are likely to be integrated in broadband ATM networks, and suggest schemes for bandwidth allocation and transmission scheduling to meet the quality and performance objectives. We propose ATM cell-multiplexing using a Dynamic Time-Slice (DTS) scheme which guarantees a required bandwidth for each traffic class and/or virtual circuit (VC), and is dynamic in that it allows the different traffic classes or VCs to share the bandwidth with a soft boundary. Any bandwidth momentarily unused by a class or a VC is made available to the other traffic present in the multiplexer. The scheme guarantees a desired bandwidth to connections which require a fixed wide bandwidth. Thus, it facilitates setting up circuit-like connections in a network using the ATM protocol for transport. The DTS scheme is an efficient way of combining constant bit-rate (CBR) services with variable bit-rate (VBR) statistically multiplexed services. We also describe methodologies to schedule delivery of delay-tolerant data traffic within the framework of the DTS scheme. Important issues such as buffer allocations, guarantee of service quality, and ease of implementation are also discussed.

1. INTRODUCTION

Broadband/high-speed packet networks based on the Asynchronous Transfer Mode (ATM) standard are currently being designed for integration of various traffic types, ranging from voice and interactive data to image, bulk data and video [1] [2]. These call types vary widely in their bandwidth requirements and tolerance to message delay and loss. There is a critical need for developing methodologies to efficiently integrate these traffic types on a single high-speed packet network. The challenge is to develop methods that meet the performance objectives for different traffic types, while making the best possible use of all the network resources, i.e., transmission bandwidth, buffers, and packet-switch capacity.

Significant work on packetized voice/data integration at lower speeds (DS1 at 1.536 Mb/s and E1 at 2.048 Mb/s) was done recently as part of the AT&T Wideband Packet Technology (WPT) project. The WPT technology is implemented in a system known as the Integrated Access and Cross-connect System (IACS) [3] [4]. The IACS integrates packetized voice, data, and signaling using efficient bandwidth allocation and congestion control methods[4]. Different congestion control methods seem to be well suited for different traffic types. For example, congestion control can be done on voice traffic by bit- or cell-dropping [5] [6]. while it is done on the data traffic by combination of methods such as window-based flow controls, bandwidth usage monitoring, marking cells (or packets), discarding the marked cells during congestion, rate-based flow controls, and various combinations of these schemes [7] [8] [9] [10] [11] [12] [13] [14]. Continued efforts are necessary to investigate new congestion avoidance and control methods for integration of voice, image, video, and high speed data in broadband networks. It is possible to achieve high efficiency of statistical multiplexing for narrowband traffic, e.g., voice and interactive data, which tend to obey the law of large numbers and have low burstiness when many sources are superposed together in a high-speed packet/cell multiplexer [5]. Isochronous high bandwidth services, e.g., fixed bit-rate video, would require guaranteed bandwidth. Medium range bandwidth services, such as variable bit-rate video, image, file transfers, and bulk data, can be extremely bursty [15], and call for new studies to characterize traffic and evaluate the statistical multiplexing efficiency.

A traffic classification approach for high-speed networks is presented in Section 2. The DTS scheme and some servicing strategies based on DTS are described in Section 3. In Section 4, we present a queueing analysis for prediction of mean and standard deviation of delay, and capacity estimation for a statistical multiplexer serving high speed bursty sources. The summary and conclusions are provided in Section 5.

2. TRAFFIC CLASSIFICATION IN BROADBAND/HIGH-SPEED NETWORKS

A variety of applications with widely varying bandwidth and delay requirements will be supported in an integrated fashion in broadband ATM networks (see Figure 1). For this reason, it is very helpful to categorize traffic at the applications layer as follows:

Type 1A: Delay-sensitive Isochronous High Bandwidth Services

This traffic type is one of two types (1A and 1B) of real-time high bandwidth services (RT-HBS). It is isochronous because it requires a fixed high bandwidth for the duration of a call. It requires real-time service, and therefore, guaranteeing bandwidth at call set-up time is necessary. Examples of this traffic are constant bit-rate (CBR) conference video and other real-time high-bandwidth CBR services.

Type 1B: Delay-sensitive Non-Isochronous High Bandwidth Services

This traffic is also RT-HBS type but it is non-isochronous in that each call alternates between active and inactive periods, and generates a high-bandwidth data-stream during the active periods. Examples of this type of traffic are variable bit-rate (VBR) video and LAN-to-LAN interconnects. This traffic is further divided into sub-classes. For example, VBR video and LAN-to-LAN calls will be separately grouped for admission control, bandwidth allocation, and buffer management.

Type 2: Delay-insensitive High Bandwidth Services

These are essentially non-real-time bulk information transport services. The user may specify the extent of delay that is tolerable; specified in the units of minutes, hours, or as over-night delivery. There are different alternative overlay methods for appropriately serving this traffic type as will be discussed later. Essentially, the network transports this type of traffic during slack periods. Hence, it calls for suitable ways of storing the bulk information, identifying network slack periods, and scheduling of the transmission. Examples of this traffic type are: document/image/video delivery services (non-real-time, sometimes overnight delivery).

Type 3: Low-bandwidth Statistically Multiplexed Services (SMS)

Calls of this type are delay sensitive with end-to-end delay requirements ranging from a few tens of milliseconds to a few hundreds of milliseconds. The examples of such traffic are packetized voice, inter-active data, enquiry-response messages, etc. No individual bandwidth reservations are made on a per call basis for calls of this type. However, the network will reserve a certain bandwidth for this type of traffic collectively. There will be randomness in the packet arrival process due to activity/inactivity alternations at the source. The number of calls of this traffic type will be restricted at each node so that the anticipated average utilization of bandwidth by this traffic (collectively) does not exceed a certain predetermined fraction (say, 60 or 70%) of the bandwidth guaranteed for it.

3. BANDWIDTH ALLOCATION AND CONGESTION AVOIDANCE IN B-ISDN

The detailed traffic classification (presented in Section 2) is needed to implement the resource sharing methods described in this section. We present a specific cyclic service discipline, called Dynamic Time-Slice (DTS), for achieving bandwidth efficiency, congestion avoidance, and fairness across all traffic classes.

3.1 THE DYNAMIC TIME-SLICE (DTS) SCHEME

The DTS scheme, which is a generalization of the (T_1, T_2) - scheme (described in [4]) to multiple traffic classes, can be explained with the help of Figure 2. Each isochronous HBS call (Type 1A) is serviced through a separate queue. Likewise, each non-isochronous traffic class (Type 1B) is also assigned a queue of its own. Accordingly, in Figure 2, variable bit-rate (VBR) video and LAN-to-LAN traffic are statistically multiplexed in separate queues. Similarly, the low-bandwidth SMS traffic classes (Type 3), such as voice and low-speed data, are also assigned separate queues. And each delay-insensitive delivery (Type 2) traffic is also assigned a separate queue when it is scheduled for service. The two lowest queues in Figure 2 are shown servicing a bulk data request and a video delivery request. In Figure 2, the lower most two queues are set up to serve Type 2 calls because there is spare bandwidth on the link.

The DTS server shown in Figure 2 services all the queues by cyclically visiting each queue and allocating a slice of time to it. The time slice allocated to a queue would be proportional to the bandwidth required by that queue. For example, in Figure 2, 30 cells in a cycle of 450 cells are allocated to the voice queue to give 10 Mb/s ($\frac{1}{15}$ th of the link bandwidth) to all the voice calls collectively. Let M_c denote the maximum number of cells that can be transmitted in a DTS cycle. For the example in Figure 2, M_c = 450. The number, M_c, should be large enough so as to facilitate a choice of a wide range of fractions of link bandwidth; a larger value of M_c means a finer granularity for bandwidth allocations. However, M_c should be chosen so that the cycle length remains of the order of a few milliseconds (e.g., 1 or 2 ms). This will guarantee that the maximum delay for the cells of Type 1A calls will be limited to a millisecond or two.

Let n denote the number of queues in the DTS configuration at a particular instance. At the instance pictured in Figure 2, n is 8. In general, n changes as the number of traffic classes that offer calls varies (on the link in consideration). Also, n varies as the number of Type 1A or Type 2 calls (requiring dedicated queues) vary. As n varies, the time-slice allocations, $T_1, T_2, \ldots, T_n$, for the n queues, are also reassigned values to reflect the proportionate bandwidth requirements for all traffic classes or HBS calls in service. The units of the time-slices are in terms of the number of ATM cells (see Figure 2). We will shortly describe how the values of $T_1, T_2, \ldots, T_n$ are selected.

Though not illustrated in Figure 2, we assume that there is one queue at each link that is dedicated to the signaling traffic; we call it queue 0. This queue would have a generous allocation of buffers to ensure no cell loss (or, probability of loss < 10^{-9}), and its time-slice parameter T_0 would be chosen so that signaling cells would have negligible probability of waiting more than one DTS cycle.

Computation of DTS Time-Slice Parameters:

Suppose that $n+1$ queues are to be set up with bandwidth requirements in the fractions of $f_0, f_1, \ldots, f_n$ of the link bandwidth. The DTS time parameters $(T_0, T_1, \ldots, T_n)$ are chosen such that the following relationships hold:

$$f_i = \frac{T_i}{\sum_{i=o}^{n} T_i} \quad , \quad 0 \le i \le n \ , \tag{1}$$

$$\sum_{i=1}^{n} f_i \le 1 - f_0 \ , \tag{2}$$

$$\sum_{i=0}^{n} T_i \le D_c \, ms \qquad (D_c \approx 1 \ to \ 2 \ \ ms) \ , \tag{3}$$

where $D_c = M_c \tau$ is the DTS service cycle time, M_c is the number of cells in a DTS service cycle as defined before, and τ is the cell transmission time on the link. Eq. (1) ensures that a fraction f_i of the link bandwidth is guaranteed for the queue i. Next, (2) indicates that the total allocated bandwidth to all traffic classes can not exceed a fraction $(1-f_0)$ of the link capacity. This is because a fraction f_0 is assumed reserved for the signaling traffic on each link. Finally, (3) indicates that the DTS service cycle time should be about one to two milliseconds. This will guarantee that the cell delays never exceed D_c ms for all Type 1A (isochronous HBS) calls. Thus the scheme guarantees not only a bandwidth but also a low maximum delay for the cells of Type 1A calls. This enables providing constant bit-rate (CBR) service with a low end-to-end fixed delay. The queueing delays experienced by cells of other statistically multiplexed traffic classes (Types 1B, 3) will also be typically no more than a few milliseconds because the call acceptance rules will ensure moderate occupancy in their queues. When call admissions/departures occur, the DTS parameters can be easily recomputed within a matter of a few microseconds so that (1)-(3) are still satisfied.

3.2 DETAILS OF SERVICING STRATEGY

A variety of servicing schemes are applied for call-admittance, and set-up, and they depend on the characteristics and requirements of the various traffic classes. We now describe these schemes.

Servicing scheme for isochronous HBS traffic (Type 1A):

- User makes a call setup request for allocation of a virtually reserved peak bandwidth,
- Network sends a bandwidth allocation request to all nodes enroute to the destination,
- The user is guaranteed the allocated bandwidth for the duration of the call,
- An efficient routing scheme is used to find an appropriate path for bandwidth reservation.

Servicing scheme for non-isochronous HBS traffic (Type 1B):

- Calls within each call-class are statistically multiplexed but each call-class (e.g., LAN-to-LAN, VBR video, etc.) is given a guaranteed bandwidth,
- Bandwidth vs. capacity tables (i.e., traffic tables) are maintained for each call-class. A traffic table provides the amount of bandwidth that is needed to multiplex a given number of calls in the queue for that call-class. It is determined by taking into consideration the burstiness, coding type, and performance requirements for the call-class in consideration. (See Section 4 for details.)
- If spare bandwidth is available on the link, a new call is admitted and the bandwidth (i.e., time-slice allocation) for this call is increased as specified in the capacity vs. bandwidth tables,
- When a call completes, the bandwidth is relinquished by lowering the DTS time-slice value for its class,
- Placing each call-class in a separate queue, with a virtual bandwidth pipe, facilitates each call-class being served according to its performance requirements,
- Any unused bandwidth, due to statistical fluctuations in the activity of calls in progress, is momentarily made available to other traffic in the system,
- An efficient routing scheme is used to find an appropriate path for each call,
- The capacity vs. bandwidth tables can be refined or updated as traffic characteristics become better known, or technology improves (e.g., use of more efficient video compression methods).

We also describe three schemes (A, B, and C) for servicing delay-insensitive HBS traffic (Type 2). Servicing schemes A and B differ depending on whether the information storage is located in the network edge node or in the CPE (for the requests that are awaiting scheduled delivery). While schemes A and B are based on the notion of deferred delivery with bandwidth reservation, scheme C transmits the cells of this class with low priority and no bandwidth guarantee. Thus, in scheme C the cells belonging to this class use any cell-slots that are unused by all other classes in the DTS service cycles.

Servicing-scheme A for delay-insensitive HBS traffic (Type 2):

- User sets up a connection with the nearest service node (not end-to-end),
- Service node accepts the bulk information and stores it in bulk memory,
- User disconnects after notifying the intended destination to the service node,
- The second stage of call set up is done by the network during slack periods of traffic,
- The service node negotiates with other nodes enroute to designated destination for an appropriate bandwidth guarantee (so that the bulk information transmission can be completed within the user-specified delivery dead-line),
- Now the bulk data is transported across,
- Service node at receiving end receives and stores the bulk data, and notifies the intended end user,
- Call is cleared when the data is delivered,
- Sender is sent a service completion message.

Servicing-scheme B for Type 2 traffic:

- User specifies the size of transaction, delay tolerance, and bandwidth required,
- Network registers a request for call setup,
- The bulk information awaits transmission at the user terminal rather than in the access node,
- Call is setup during network slack period but within deadline specified by user,
- The bulk information is transported with an appropriate bandwidth guarantee (see description in scheme A above).

Servicing-scheme C for Type 2 traffic:

- A low priority queue is established at each link,
- This queue is low priority in the sense that it has no bandwidth guarantee; it utilizes spare bandwidth (idle cell-slots) in each DTS cycle,
- The Type 2 traffic is served through this queue,
- In order to assure cell loss due to buffer overflow, a STOP/SEND signaling scheme is used (backward congestion notification),
- Based on a threshold on the low-priority buffer fill, each node sends STOP/SEND signals to neighbors (see Section 5 of [22] for a related discussion),
- The queued low-priority traffic at a link is transmitted when spare bandwidth is available and the link has SEND status.

The network can have the capability to provision all three schemes (A, B, and C) described above, and the choice can be based on the needs specified by individual Type 2 call requests. Finally, the following is a description of the servicing strategy for Type 3 calls.

Servicing scheme for low- or medium-bandwidth SMS traffic (Type 3):

- Similar to the non-isochronous HBS calls (Type 1B), but the call setup procedures (acceptance, routing) are likely to be different,
- Each Type 3 call class, e.g., voice, low-speed data, interactive image, will have its own queue,
- The law of large numbers holds better for this traffic and the statistical multiplexing efficiency can be much higher compared to Type 1B,

4. BURSTINESS, STATISTICAL EFFICIENCY, AND CAPACITY

The DTS scheme guarantees a bandwidth for each call-class; hence the call-admittance procedures can be determined individually for each call-class. The bandwidth vs. capacity tables (i.e., the traffic tables), mentioned earlier in Section 3.2, can be computed separately for each call-class assuming that it has an exclusive bandwidth allocation. This approach results in a somewhat conservative but tractable allocation of bandwidth (and call-admittance), because the multiple call-classes actually share the link bandwidth with soft boundaries (see description of the DTS scheme in Section 3). The computation of a traffic table will depend on a careful analysis and quantification of traffic burstiness, statistical multiplexing efficiency, and capacity for the call-class in consideration. These issues are well understood for packetized voice[5][6] [16]. There is significant work in progress on modeling variable bit-rate video traffic in ATM networks [15] [17]. Several other studies also report on other specific models to capture the effects of burstiness for various specific traffic types and stochastic characterizations (for example, see [18] [19] [20]. The actual traffic capacity and performance (e.g., latency) results would depend on accurate modeling and realistic traffic assumptions for each call-class.

4.1 MODELING AND CAPACITY PREDICTION

Here we present an analysis of, and some insights into, the relationship between burstiness and capacity for high-bandwidth statistically varying cell streams. Each call-class has a guaranteed minimum bandwidth of its own in the DTS scheme. Our study assumes that only this minimum bandwidth is available to each call-class in isolation, and therefore the results we obtain are conservative.

To facilitate the discussion, let us assume that each source alternates between bursty and idle periods (see Figures 3 and 4). When the bursts from multiple sources overlap, the instantaneous total cell arrival rate can exceed the cell service rate on the link. In order to meet the cell delay (and cell loss) objectives, the duration of such overload periods should be typically small compared to the cell delay objectives (and much smaller than the buffer size). The statistical multiplexing efficiency or gain, G, is the ratio of the maximum number of sources that can be statistically multiplexed (on the link while meeting the delay/loss objectives) to the number of sources that can be deterministically multiplexed by allocating the peak bandwidth to each source. As illustrated in Figure 3, the statistical multiplexer gain, G, depends on various factors such as the average burst duration, the source peak rate, and the link bandwidth, etc. We now quantify the effects described here by means of a queueing model.

In Figure 4, we illustrate a bursty source with peak rate R Kbps, and mean active and inactive periods, z_1 and z_2, which are specified in terms of multiples of the cell inter-arrival (i.e., cell formation) time, T, of the source (see (5) below). The source activity factor, a, is given by:

$$a = \frac{z_1}{z_1 + z_2} . \tag{4}$$

Let B denote the guaranteed link bandwidth for the traffic class in consideration. Let X_P denote the octets in an ATM cell that contain actual information bits from a source. Let H denote the total overhead octets including the ATM header and the adaptation layer control overhead. Then the cell inter-arrival time, T (during the burst periods of a source), and the cell transmission time, τ, on the link are given by:

$$T = \frac{8\,X_P}{R} \quad , \quad \tau = \frac{8\,(X_P + H)}{B} . \tag{5}$$

Let $N_{\min}$ denote the number of sources that can be multiplexed using the peak rate allocation. Let $N_{\max}$ denote the maximum number of sources that can be multiplexed (theoretically) for a given link bandwidth, B, and activity factor, a. Then, $N_{\min}$ and $N_{\max}$, and the corresponding bounds, $G_{\min}$ and $G_{\max}$, on the statistical multiplexer gain, G, are given by:

$$N_{\min} = \frac{T}{\tau} = \frac{X_P}{X_P + H}\,\frac{B}{R} , \quad N_{\max} = \frac{N_{\min}}{a} = \frac{X_P}{X_P + H}\,\frac{B}{a\,R} , \tag{6a}$$

$$G_{\min} = 1 , \quad G_{\max} \equiv \frac{N_{\max}}{N_{\min}} = \frac{1}{a} . \tag{6b}$$

The burstiness of a single source can be represented by the squared-coefficient of variation, c_1^2, of its inter-arrival time, U. For the source shown in Figure 4, let us assume that the burst lengths are geometrically distributed with mean z_1, and the inactive periods are exponentially distributed with mean z_2. Now it can be shown that c_1^2 is given by[16]:

$$c_1^2 \equiv \frac{var\,(U)}{[E(U)]^2} = (1-a)^2\,(2z_1 - 1) . \tag{7}$$

Eq. (7) can be derived from (2) in [16] by a change of notation and with some algebraic manipulation (note that p in [16] is $(z_1 - 1)/z_1$ in the present notation, and β in [16] is $1/z_2T$).

For the modeling of the aggregate cell arrival process resulting from superposition of bursty sources, we follow a two-parameter approximation similar to that given in [16] [21]. Essentially, the correlations in the superposition process are captured in the queueing analysis by appropriately inflating (at high loads) the effective squared coefficient of variation, c_a^2, of the aggregate arrival process. Let n denote the number of sources multiplexed, and ρ denote the corresponding link load. Then we have,

$$\rho = \frac{a\,R\,n}{B} . \tag{8}$$

The following sequence of equations provide an expression for the mean cell delay, $E[W]$:

$$\omega = \frac{1}{1 + 4(1-\rho)^2(n-1)} , \tag{9}$$

$$c_a^2(\rho,n) \equiv c_a^2 = \omega\,c_1^2 + (1-\omega) , \tag{10}$$

$$c_m^2(\rho,n) \equiv c_m^2 = \begin{cases} (1 - \frac{\rho}{a}) + \frac{\rho}{a}\,c_a^2 , & \text{for } 0 \le \rho < a , \\ c_a^2 , & \text{for } \rho \ge a , \end{cases} \tag{11}$$

$$E[W] = \frac{c_m^2\,\rho\,\tau}{2(1-\rho)} . \tag{12}$$

In (11) above, we use the source activity factor, a, as a transition point for gradual transition from the Poisson approximation to the two-parameter approximation. This is because at low loads (i.e., $\rho \le a$), the simple Poisson approximation is fairly accurate (see [6][16]). The burstiness of the sources is reflected in the superposition process and in-turn in the queue performance at higher loads only (i.e., $\rho > a$). To obtain the standard deviation, σ_w, of the cell delays, we have the following sequence of equations based on an approximation given in [21]:

$$h = \frac{4\rho}{(1 + 4\rho^2)\,c_m^2} , \tag{13}$$

$$\sigma \equiv Prob.(W > 0) = \rho + (c_m^2 - 1)\,\rho\,(1-\rho)h , \tag{14}$$

$$c_d^2 = \frac{2\rho + 1}{3} , \tag{15}$$

$$c_w^2 = \frac{c_d^2 + 1 - \sigma}{\sigma} , \tag{16}$$

$$\sigma_w = c_w\,E\,[W] . \tag{17}$$

In the above equations, D denotes the conditional delay given that the server is busy, σ denotes the probability of delay (i.e., $\sigma = Pr[D > 0]$), c_d^2 denotes the squared coefficient of variation of D, and c_w^2 denotes the squared coefficient of variation of the waiting time W. For details of derivations and validation see [16][21]. The equations above are a somewhat modified version of those in [21] because the service time in the present system is deterministic.

We use D_b as defined below to be a performance measure of interest:

$$D_b \equiv E[W] + 5\,\sigma_w = (1 + 5\,c_w)\,E[W]\,. \qquad (18)$$

The cell delay should rarely exceed D_b for many common traffic types, provided the load on the system is at least moderately below the saturation level. If the mean and the standard deviation of delay are estimated fairly accurately (which we believe is the case for the model presented here[16][21]), then D_b is a reasonable estimate for the buffer size to achieve negligible cell overflow probability. Alternatively, requiring D_b to be less than a prespecified buffer size could be used as a performance measure for capacity estimation (see examples below).

4.2 NUMERICAL EXAMPLES AND DISCUSSION

In the numerical examples illustrated in Figures 5-8, the source activity is assumed to be 40% (i.e., $a = 0.4$). The source peak rate, R, is assumed to be 3 Mb/s except for the case of the 1.5 Mb/s link when it is assumed to be 64 Kb/s. The actual number of information octets in an ATM cell is assumed to be 44 (i.e., X_P = 44) allowing five octets for the ATM header (see (5)) and four octets for the adaptation layer control.

In Figure 5, the plots of D_b are shown as a function of the percent link load for DS3 (45 Mb/s) and STS-3c (150 Mb/s) link speeds. The burstiness or c_1^2 value of the cell arrival process (from a source) is expected to vary over several orders of magnitude for different classes of traffic in high speed networks. Hence, in our numerical examples, we varied c_1^2 over a wide range of values. In Figure 6, the plots of D_b are shown as a function of the load in terms of the number of calls multiplexed on the link.

It is a desirable objective in high-speed networks to limit the queueing delay at individual nodes to be less than a few milliseconds. We use a value of 2 milliseconds as an objective for D_b to determine the statistical multiplexer gain, G. The plots of G are shown in Figure 7 for various link speeds as a function of the burstiness (measured by c_1^2). The plots of the capacity (in terms of the number of sources) as a function of burstiness are shown in Figure 8. The burstiness measure, c_1^2, is a function of the burst length, z_1, and the activity factor, a (see (7)). Thus, the mean cell delay tends to be proportional to the length of the burst or active period (see (7),(10)-(12)). The burst length would also have a significant impact on the cell loss probability for finite buffer capacity. The delay and the cell loss probability will determine the call capacity of a multiplexer for a given bandwidth. When more bandwidth is available, the burstiness effects are mitigated somewhat due to the interaction of a larger number of sources, and the capacity is correspondingly higher for higher bandwidth links (see Figures 7, 8).

As new coding and compression methods are developed for voice, image, and video, the traffic capacity tables for these classes will be updated accordingly. By way of this updating capability, the DTS servicing strategy adapts to changes in technology and refinements in modeling tools for traffic characterization and capacity prediction.

5. SUMMARY AND CONCLUSIONS

The main features of the Dynamic Time-Slice (DTS) scheme can be summarized as follows:

- The DTS scheme guarantees a desired bandwidth to connections which require a fixed wide bandwidth. Thus it facilitates setting up circuit-like (i.e., CBR) connections in an ATM network.
- Statistical multiplexing is done within each call-class when possible. Call-admittance is based on capacity vs. bandwidth tables for that class. These traffic tables are updated as technology (e.g., video coding/compression) changes.
- The DTS scheme is seen as an efficient way of combining isochronous (i.e., constant bit-rate or CBR) High Bandwidth Services (HBS) with variable bit-rate (VBR) statistically multiplexed connections.
- Placing each call-class in a separate queue, with a virtual bandwidth pipe, facilitates each call-class to be served according to its own performance requirements.
- Any bandwidth unused by a call-class or an isochronous HBS call is momentarily made available to other traffic present in the multiplexer.
- The delay-insensitive bulk information is stored in the CPE or service node, and transmitted via a secondary call setup during network slack periods. This is an overlay technique for scheduling and resource management, and can be used with ATM or any other protocol. We also described other overlay techniques (see Section 3.2) for bulk information (i.e., delay-insensitive HBS) services.
- Users of HBS service may be penalized by buffer overflow if they exceed the contracted bandwidth when there is no spare (or momentarily unused) bandwidth on the link.
- The call admittance procedures are call-class based, and hence the network resource usage costs can be estimated separately for each call-class. Thus the DTS framework could facilitate a simple call-class based billing.

In this paper, we also identified ways of quantifying the bandwidth-burstiness-capacity relations. There is still a critical need to develop network routing strategies to find efficient routes for bandwidth reservation and call setup. Also, traffic measurements are needed for continued characterization of new traffic types and performance predictions. In the future, it is important that bandwidth allocation and congestion control methods are evaluated in conjunction with suitable billing methods. A lead in this direction is provided by the call-class based resource management (via the DTS scheme) presented in this paper.

A more detailed report on this work is available elsewhere (please see Sriram[22]).

6. ACKNOWLEDGEMENTS

The author wishes to thank many colleagues at the AT&T Bell Laboratories in New Jersey and the AT&T-NSI in the Netherlands for helpful discussions and comments related this work.

REFERENCES

1. *IEEE Journal on Selected Areas in Communications*: Special issue on "Architectures and Protocols for Integrated Broadband Switching," December 1991.

2. *Computer Networks and ISDN Systems Journal* : Special issue on "The ATM-Asynchronous Transfer Mode", Vol. 24, No. 4, May 1992.

3. M.H. Sherif, A.D. Malaret-Collazo, and M.C. Gruensfelder, "Wideband Packet Technology in the Integrated Access and Cross-Connect System (IACS)," *International Journal of Satellite Communications*, vol. 8, pp.437-444 (1990).

4. K. Sriram, "Dynamic Bandwidth Allocation and Congestion Control Schemes for Voice and Data Integration in Wideband Packet Technology," *Proc. of ICC'90*, Atlanta, GA, pp. 1003-1009.

5. K. Sriram, R.S. McKinney, and M.H. Sherif, "Voice Packetization and Compression in Broadband ATM Networks," the *IEEE Journal on Selected Areas in Communications*, April 1991, pp. 294-304.

6. K. Sriram and D. M. Lucantoni, "Traffic Smoothing Effects of Bit Dropping in A Packet Voice Multiplexer," the *IEEE Trans. on Commun.*, July 1989, pp. 703-712.

7. J.Y. Hui, M.B. Gursoy, N. Moayeri, and R. D. Yates, "A layered broadband switching architecture with physical or virtual path configurations," *IEEE J. Select. Areas Commun.*, December 1991, pp. 1416-1426.

8. A.E. Eckberg, B.T. Doshi, and R. Zoccolillo, "Controlling congestion in B-ISDN/ATM: Issues and Strategies," *IEEE Communications Magazine*, September 1991, pp. 64-74.

9. G.M. Woodruff and R. Kositpaiboon, "Multimedia traffic management principles for guaranteed ATM network performance," *IEEE J. Select. Areas Commun.*, vol. 8, April 1990, pp. 437-446.

10. S.J. Golestani, "A framing strategy for congestion management," *IEEE J. Select. Areas Commun.*, September 1991, pp. 1064-1077.

11. J. M. Hyman, A. A. Lazar, and G. Pacifici, "Real time scheduling with quality of service constraints," *IEEE Journal on Selected Areas in Communications*, September 1991, pp. 1052-1063.

12. C. R. Kalmanek, H. Kanakia, and S. Keshav, "Rate Controlled Servers for Very High Speed Network," *Globecom'90*, pp. 300.3.1-9.

13. G. Ramamurthy and R.S. Dighe, "A multidimensional framework for congestion control in B-ISDN," *IEEE J. Select. Areas Commun.*, December 1991, pp. 1440-1451.

14. P.E. Boyer and D.P. Tranchier, "A reservation principle with applications to the ATM traffic control," *Computer Networks and ISDN Systems Journal*, Vol. 24, No.4, May 1992, pp. 321-334.

15. D. Heyman, A. Tabatabai, and T.V. Lakshman, "Statistical analysis and simulation study of video teleconferencing traffic in ATM networks," *Proc. of IEEE Globecom'91*, pp. 21-27.

16. K. Sriram and W. Whitt, "Characterizing superposition arrival processes in packet multiplexers for voice and data," *IEEE Journal on Selected Areas in Commun.*, vol. SAC-4, no. 6, September 1986, pp. 833-846.

17. A.R. Reibman and A.W. Berger, "Average-Rate Traffic Descriptors for VBR Video Teleconferencing over ATM networks," *Globecom'92*.

18. K.S. Meier-Hellstern, P.E. Wirth, Y-L. Yan, and D.A. Hoeflin, "Traffic models for ISDN users: Office automation application," submitted to the *IEEE Trans. on Commun.*

19. I. Norros, J.W. Roberts, A. Simonian, and J. Virtamo, "The superposition of variable bit rate sources in an ATM multiplexer," the *IEEE Journal on Selected Areas in Communications*, April 1991, pp. 378-387.

20. I. Stavrakakis and S. Tsakiridou, "A unified analysis of integrated services TDM under various policies and packet arrival processes," *Proc. of IEEE INFOCOM'92*, May 1992, Florence, Italy.

21. W. Whitt, "The queueing network analyzer," *Bell Syst. Tech. J.*, Part 1, vol. 62, no. 9, pp. 2779-2815, Nov. 1983.

22. K. Sriram, "Methodologies for Bandwidth Allocation, Transmission Scheduling, and Congestion Avoidance in Broadband ATM Networks," *Computer Networks and ISDN Systems Journal:* special issue on "Traffic Issues in ATM Networks" (to appear).

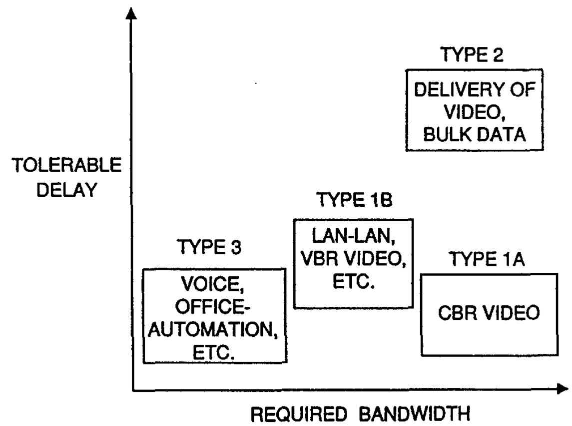

FIGURE 1. Traffic types and their delay/bandwidth requirements.

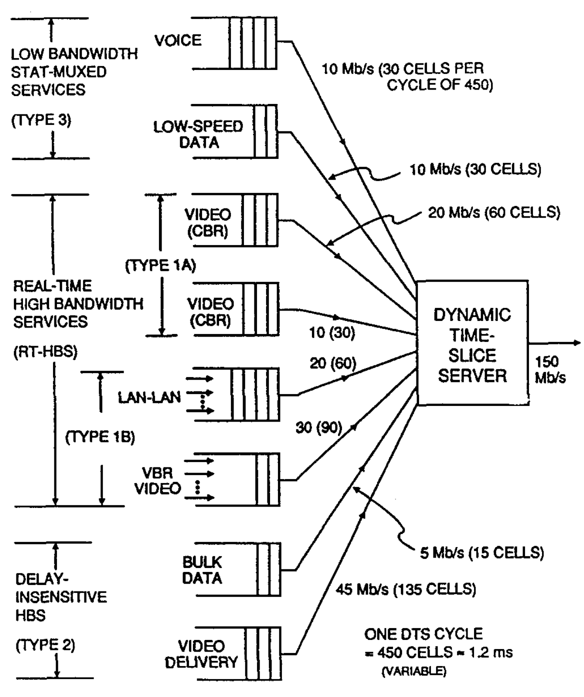

Figure 2. An illustration of operation of the DTS scheme.

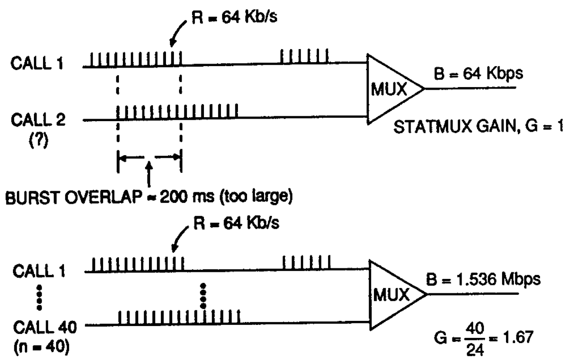

Figure 3. Illustration of relationship between burstiness and statistical multiplexing efficiency. (The call activity factor is assumed to be about 40%.)

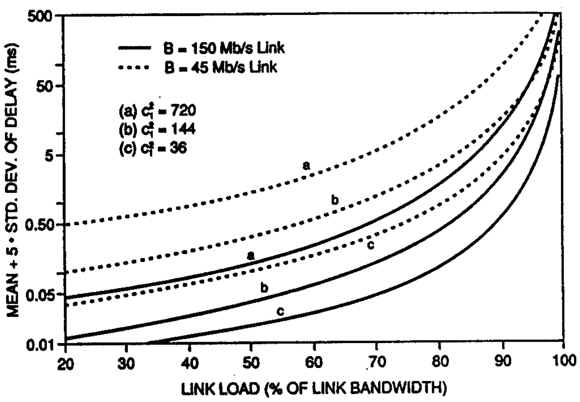

Figure 5. Mean plus five standard deviations of delay as a function of percent load on the link. The link bandwidth and the squared coefficient of variation (c_i^2) are varied.

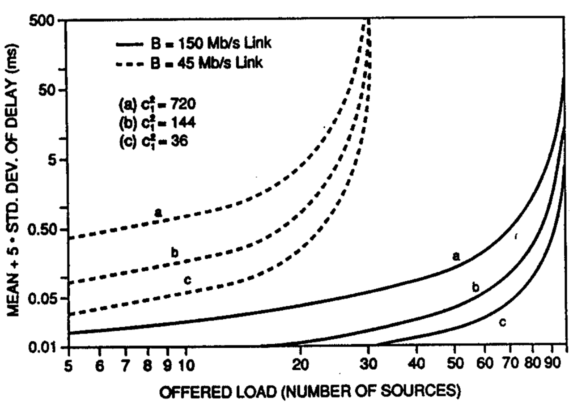

Figure 6. Mean plus five standard deviations of delay as a function of the number sources multiplexed on the link. The link bandwidth and the squared coefficient of variation (c_i^2) are varied.

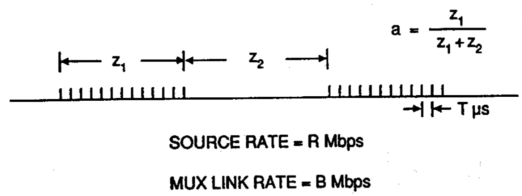

Figure 4. Illustration of bursty nature of a source.

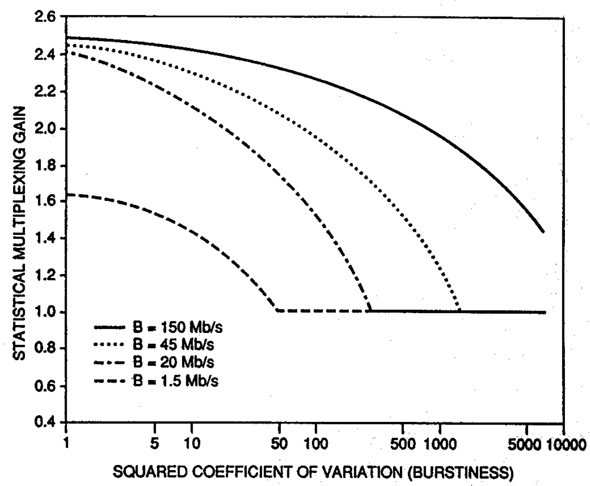

FIGURE 7. Statistical multiplexing gain as a function of the burstiness (estimated by the squared coefficient of variation of the arrival process of a single source). The gain is higher when many more sources are multiplexed together at higher link bandwidths. The source activity is assumed to be 40% and peak rate is assumed to be 3 Mb/s (except for the 1.5 Mb/s link case when it is 64 kb/s).

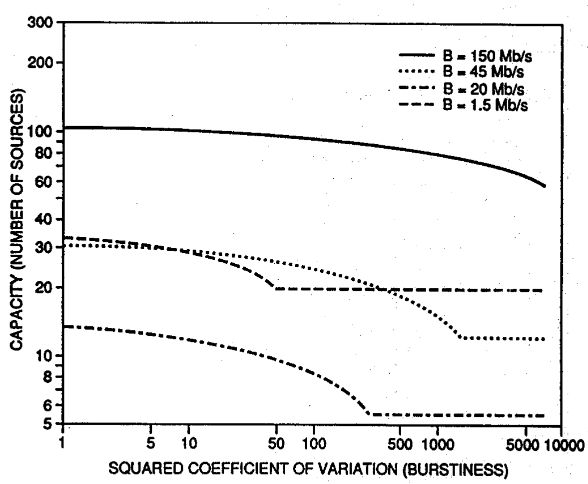

FIGURE 8. Statistical multiplexer capacity as a function of the burstiness (estimated by the squared coefficient of the arrival process of a single source). The source activity is assumed to be 40% and peak rate is assumed to be 3 Mb/s (except for the 1.5 Mb/s link when it is 64 kb/s).